

Uniform growth rate

Kasra Rafi* and Jing Tao[†]

Abstract


In an evolutionary system in which the rules of mutation are local in nature, the number of possible outcomes after m mutations is an exponential function of m but with a rate that depends only on the set of rules and not the size of the original object. We apply this principle to find a uniform upper bound for the growth rate of certain groups including the mapping class group. We also find a uniform upper bound for the growth rate of the number of homotopy classes of triangulations of an oriented surface that can be obtained from a given triangulation using m diagonal flips.

1 Introduction

Let G be a group and S be a generating set for G . We denote the word length in G associated to S with $\|\cdot\|_S$. Recall that the growth rate of G (relative to S) is defined to be

$$h_G = \lim_{R \rightarrow \infty} \frac{\log \#B_R(G)}{R}, \quad \text{where} \quad B_R(G) = \{g \in G \mid \|g\|_S \leq R\}.$$

In his 60th birthday conference, Bill Thurston mentioned that the mapping class group has a growth rate that is independent of its genus. Namely, consider the following set of curves on a surface $\Sigma = \Sigma_{g,p}$ of genus g with p punctures:



Let S be the set of Dehn (or half) twists around these curves. This set S generates $\text{MCG}(\Sigma)$, the mapping class group of Σ [Lic64, FM12, Art47, Bir74]. (Note that S is a combination of the Lickorish generators of the mapping class group of a closed surface and the standard generators of a

*Partially supported by NCERC Research Grant, RGPIN 435885.

[†]Partially supported by NSF Research Grant, DMS-1311834

braid group). We will refer to S as the set of *extended Lickorish generators* for $\text{MCG}(\Sigma)$. Then the growth rate of $\text{MCG}(\Sigma)$ equipped with the word metric associated to S has an upper bound that is independent of the topology of Σ . This, Thurston asserted, is true since most pairs of elements in S commute.

Note that, in fact, the number of elements in S that do not commute with a given element in S is uniformly bounded. We show that this is enough to obtain the uniform growth rate in general.

Theorem A. *Given any c_0 , let S be a generating set for a group G such that, for every $s \in S$, the number elements of S that do not commute with s is bounded by c_0 . Then $h_G \leq \log(2c_0 + 2) + 1$.*

Since each curve in the extended Lickorish generators intersects at most 3 other curves, we obtain:

Corollary B. *The growth rate of $\text{MCG}(\Sigma)$ relative to the extended Lickorish generators is bounded by $\log 8 + 1$.*

Uniform growth rate can also be shown regarding groups $\text{Aut}(F_n)$, $\text{Out}(F_n)$, $\text{GL}_n(\mathbb{Z})$ and similar groups if the generating set is chosen such that the number of generators that do not commute with a given generator is uniformly bounded. In fact, these groups have natural generating sets with this property. For example, in the case of $\text{Aut}(F_n)$, let F_n be the free group with basis $\{a_1, \dots, a_n\}$, and consider the following three types of automorphisms of F_n .

1. Inversion: For $1 \leq i \leq n$, $I_i(a_i) = \overline{a_i}$ and fixes all other a_j .
2. Transposition: For $1 \leq i \leq n-1$, $P_i(a_i) = a_{i+1}$ and $P_i(a_{i+1}) = a_i$ and fixes all other a_j .
3. Multiplication: For $1 \leq i \leq n-1$, $M_i(a_i) = a_i a_{i+1}$ and fixes all other a_j .

The collection of inversions, transpositions, and multiplications generate $\text{Aut}(F_n)$ [MKS66, LS77] and is called the set of *local Nielsen generators*. For each $s \in S$, the number of elements that do not commute with s is at most 7, we obtain:

Corollary C. *The growth rate of $\text{Aut}(F_n)$ relative to local Nielsen generators is bounded by $\log 16 + 1$.*

1.1 Evolving structures

Another context to apply this philosophy is the setting of evolving structures. We follow the footsteps of the work of Sleator-Tarjan-Thurston [STT92] where they showed that if a graph is allowed to evolve using a set of rule that change the graph locally, then the growth rate of the number of possible outcomes after R mutations is bounded by a constant depending on the rules of evolution and not the size of the graph. This was used in [STT92] to estimate the diameter of the space of plane triangulations equipped with the diagonal flip metric and in [RT13] to estimate the diameter of the space of cubic graphs equipped with the Whitehead move metric. Similar to their work, one can also consider the evolution of labeled graphs. Generalizing the results in [STT92] slightly, we prove:

Theorem D. *Let G be any group and Γ be a G -labeled trivalent graph (see Section 3 for definition). Let $B_R(\Gamma)$ be the set of G -labeled graphs that are obtained from Γ by at most R splits. Then,*

$$\lim_{R \rightarrow \infty} \frac{\log \#B_R(\Gamma)}{R} \leq 3 \log 4.$$

That is, the growth rate of $B_R(\Gamma)$ is independent of the size and shape of the starting graph Γ and of the group G .

As an application, we can prove a combinatorial version of Corollary B. Namely, let $\mathcal{T}_n(\Sigma)$ be the space of homotopy classes of triangulations of the surface Σ with n vertices.

Theorem E. For $T \in \mathcal{T}_n(\Sigma)$, let $B_R(T)$ be the set of triangulations in $\mathcal{T}_n(\Sigma)$ that are obtained from T using R diagonal flips. Then

$$\lim_{R \rightarrow \infty} \frac{\log \#B_R(T)}{R} \leq 3 \log 4$$

for every surface Σ and any number of vertices n .

Note that, even though Theorem E is a direct analogue of Corollary B it does not follow from it. This is because the quotient of $\mathcal{T}_n(\Sigma)$ by $\text{MCG}(\Sigma)$ has a size that goes to infinity as the number of vertices n approaches infinity.

1.2 Remarks and references

Our Theorem A follows immediately from an upper bound on the growth rate of a right-angled Artin group $A(\Theta)$ with defining graph Θ , in terms of the maximum degree of the complementary graph $\bar{\Theta}$ (Theorem 2.1). Other results relating the growth rate of $A(\Theta)$ to the shape of Θ have been obtained in the past. For instance, it was shown in [Sco07] that the growth series of $A(\Theta)$ can be computed in terms of the clique polynomial of Θ . Similar results can be found in [AP14] and [McM14]. However, the degree of $\bar{\Theta}$ cannot be recovered from the coefficients of the clique polynomial of Θ , so these results are independent from ours.

Our proof of Theorem 2.1 is related to normal forms for elements of a right-angled Artin group. A normal form for a word representing an element in $A(\Theta)$ is obtained by shuffling commuting elements and removing inverse pairs of generators of $A(\Theta)$ whenever possible ([HM95]). By fixing an ordering of $V(\Theta)$, then every element of $A(\Theta)$ admits a unique normal form, obtained by additionally shuffling lower-order letters to lower positions whenever possible. In our proof of Theorem 2.1, we construct a *canonical representative* for a given word, obtained similarly by shuffling lower-order letters to lower positions. However, we do not need to cancel inverse pairs, so the canonical representative of a word may not be in normal form.

Acknowledgements

We thank Benson Farb for useful conversations following the talk by Thurston. We thank the GEAR network for their support. We also thank the referee for helpful comments.

2 Uniform growth rates

2.1 Preliminaries

Let G be a finitely generated group. By convention, the inverse of an element $g \in G$ will be represented by \bar{g} ; and for any subset $S \subset G$, let $\bar{S} = \{\bar{s} : s \in S\}$. A *word* in $S \cup \bar{S}$ is a sequence $w = [s_1, \dots, s_R]$,

where $s_i \in S \cup \bar{S}$; R is the *length* of w . We allow the empty word whose length is 0. A word $w = [s_1, \dots, s_R]$ represents an element $g \in G$ if $g = s_1 \cdots s_R$. (The empty word represents the identity element.) By a *generating set* for G we will mean a finite set $S \subset G \setminus \{1\}$ such that every element $g \in G$ is represented by a word in $S \cup \bar{S}$. The *word length* $\|g\|_S$ of g relative to a generating set S is the length of the shortest word in $S \cup \bar{S}$ representing g . For any R , $B_R(G)$ is the set of elements of G with word length at most R . The *growth rate* (also called the entropy) of G relative to S is

$$h_G = \lim_{R \rightarrow \infty} \frac{\log \#B_R(G)}{R},$$

where the above limit exists by sub-additivity.

We remark that the growth rate of G depends on the generating set, but positivity of the growth rate does not. The growth rate of F_n relative to a basis is $\log(2n - 1)$. If G contains a subgroup isomorphic to F_2 , then h_G is strictly positive. See [GdlH97] and the references within for more details.

2.2 RAAGs

A graph is a 1-dimensional CW complex. It is *simple* if there are no self-loops or double edges.

Let Θ be a finite simple graph. Let $V(\Theta)$ and $E(\Theta)$ be the set of vertices and edges of Θ . An element of $E(\Theta)$ will be denoted by vw , where v and w are the vertices of the edge. The *complementary graph* of Θ is the graph $\bar{\Theta}$ with $V(\Theta) = V(\bar{\Theta})$ but two vertices span an edge in $\bar{\Theta}$ if and only if they do not in Θ .

The *right-angled Artin group* or RAAG associated to Θ is the group $A(\Theta)$ with the presentation:

$$A(\Theta) = \langle s_v \text{ for } v \in V(\Theta) : [s_v, s_w] = 1 \text{ for } vw \in E(\Theta) \rangle.$$

The collection $S = \{s_v\}$ will be called the *standard generating set* of $A(\Theta)$. We will often ignore the distinction between a vertex v and the generator s_v .

Theorem 2.1. *If the valence of every vertex in $\bar{\Theta}$ is bounded above by a constant c_0 , then the growth rate of $A(\Theta)$ relative to the standard generating set S is bounded by $\log(2c_0 + 2) + 1$.*

From Theorem 2.1, we derive Theorem A as corollary.

Corollary 2.2. *Let S_G be a generating set for a group G such that for every $s \in S_G$, the number elements of S_G that do not commute with s is bounded by c_0 . Then $h_G \leq \log(2c_0 + 2) + 1$.*

Proof. Let Θ be the graph with vertex set S_G and $ss' \in E(\Theta)$ if and only if $[s, s'] = 1$ in G . The natural map from $A(\Theta)$ to G taking the standard generating set S to S_G extends to a surjective homomorphism, and the hypothesis on S_G implies the valence of every vertex in $\bar{\Theta}$ is bounded by c_0 . All together, we obtain $h_G \leq h_{A(\Theta)} \leq \log(2c_0 + 2) + 1$. \square

The rest of the section is dedicated to proving Theorem 2.1.

Given a word $w = [s_1, \dots, s_R]$ in $S \cup \bar{S}$, the j -th letter of w is $w(j) = s_j$. A word $w' = [t_1, \dots, t_R]$ is a *reordering* of w if $t_1 \cdots t_R = s_1 \cdots s_R$ and there is a permutation σ such that $t_j = s_{\sigma(j)}$. We say the letter s_k in w is *ready* for position i , $i \leq k$, if s_k commutes with every s_j , for $i \leq j \leq k$.

At every vertex v of $\bar{\Theta}$, label the half-edges at v from 1 to d_v , where $d_v \leq c_0$ is the valence of v . Let n be the cardinality of $V(\bar{\Theta})$. Fix a labeling $L_0: V(\bar{\Theta}) \rightarrow \mathbb{N}$ whose image is $\{1, \dots, n\}$.

Fix $w_0 = [s_1, \dots, s_R]$. We will inductively construct a sequence w_1, \dots, w_R of words that reorders w_0 , in conjunction with a sequence L_1, \dots, L_R of labeling of $V(\bar{\Theta})$. The final word w_R will be called the *canonical representative* of $g = s_1 \cdots s_R$ induced by w_0 . (w_R depends on w_0). Along this process, we produce an encoding of the canonical representative by a sequence of integers ℓ_1, \dots, ℓ_R .

Suppose for $i \geq 0$, $w_i = [u_1, \dots, u_i, t_{i+1}, \dots, t_R]$, a labeling L_i of $V(\bar{\Theta})$, and a sequence ℓ_1, \dots, ℓ_i are given. Among $\{t_{i+1}, \dots, t_R\}$, let U be the subset of letters that are ready for position $i+1$. Pick $t \in U$ such that $L_i(t)$ is minimal among all elements of U . Set

$$w_{i+1} = [u_1, \dots, u_i, t, t_{i+1}, \dots, \hat{t}, \dots, t_R],$$

$u_{i+1} = t$. If $t \in S$, then let $\ell_{i+1} = L_i(t)$; if $t \in \bar{S}$, then let $\ell_{i+1} = -L_i(t)$. The word w_{i+1} is a reordering of w_i and hence of w_0 by induction.

We now define the labeling $L_{i+1}: V(\bar{\Theta}) \rightarrow \mathbb{N}$. Let n_i be the largest value of L_i . Let (e_1, \dots, e_d) be the half-edges of $\bar{\Theta}$ incident at t listed in order, where d is the valence of t . For each e_k , let v_k be the vertex connected to t by the edge associated to e_k . We set $L_{i+1}(v_k) = n_i + k$ for each $k = 1, \dots, d$, and $L_{i+1}(v) = L_i(v)$ for all other $v \in V(\bar{\Theta})$.

Lemma 2.3. *Let $n = \#V(\bar{\Theta})$. Then*

$$1 \leq |\ell_1| \leq |\ell_2| \leq \dots \leq |\ell_R| \leq n + c_0 R.$$

Proof. For each $i \geq 1$, we show $|\ell_i| \leq |\ell_{i+1}|$. Let $w_R = [u_1, \dots, u_R]$. We have $|\ell_i| = L_i(u_i)$. For $v \in V(\bar{\Theta})$, $L_i(v) = L_{i+1}(v)$ unless v is in the link of u_i ; in the latter case, $L_{i+1}(v)$ is bigger than the maximal value of L_i . If u_i and u_{i+1} do not commute, then u_{i+1} is in the link of u_i , therefore $L_{i+1}(u_{i+1})$ exceeds the maximal value of L_i , and in particular $L_{i+1}(u_{i+1}) > L_i(u_i)$. If u_i and u_{i+1} commute, then they were both ready for position i . In this case, $L_i(u_{i+1}) = L_{i+1}(u_{i+1})$, and u_i was chosen precisely because its label $L_i(u_i)$ is minimal among all elements in the set u_{i+1}, \dots, u_R that were ready for position i . We conclude $|\ell_i| \leq |\ell_{i+1}|$.

The largest value of L_{i+1} is at most c_0 plus the largest value of L_i . Hence the largest value of L_R is at most $n + c_0 R$. This bounds all $|\ell_i|$. \square

Set $C_R = n + c_0 R$. Let $D_R = \{\pm 1, \pm 2, \dots, \pm C_R, C_R + 1\}$ and let

$$W_R = \{(\ell_1, \dots, \ell_R) : \ell_i \in D_R \text{ and } |\ell_1| \leq \dots \leq |\ell_R|\}.$$

Proposition 2.4. *There exists an embedding of $B_R(G)$ into W_R , hence $\#B_R(G) \leq \#W_R$.*

Proof. Let $g \in B_R(G)$ have $\|g\|_S = r$. Pick any word $w = [s_1, \dots, s_r]$ representing g and let w_r be the canonical representative of g induced from w . Let (ℓ_1, \dots, ℓ_r) be the code of w_r . If $r < R$, then extend the sequence to $(\ell_1, \dots, \ell_r, \ell_{r+1}, \dots, \ell_R)$ by setting $\ell_{r+i} = C_R + 1$ for all $i = 1, \dots, R - r$. By Lemma 2.3, $(\ell_1, \dots, \ell_R) \in W_R$. This gives a map $B_R(G) \rightarrow W_R$.

To see this is an embedding, we show how to recover w_r and hence g from the sequence (ℓ_1, \dots, ℓ_R) . Recall $\bar{\Theta}$ is equipped with a cyclic ordering of the half-edges at every vertex and a labeling L_0 of the vertices from 1 to n . Let w_0 be the empty word. Suppose for $0 \leq i \leq r-1$, $L_i: V(\bar{\Theta}) \rightarrow \mathbb{N}$ and a word

$w_i = [u_1, \dots, u_i]$ are defined. If $\ell_{i+1} = C_R + 1$, then set $u_{i+1} = u_{i+2} = \dots u_R = 1$. Otherwise, Let v be the unique vertex in $\bar{\Theta}$ with label $|\ell_{i+1}| = L_i(v)$. Set $u_{i+1} = v$ if ℓ_{i+1} is positive and $u_{i+1} = \bar{v}$ if ℓ_{i+1} is negative. Let (v_1, \dots, v_d) be the vertices in the link of u_{i+1} listed in cyclic order. Let n_i be the largest value of L_i . Construct $L_{i+1}: V(\bar{\Theta}) \rightarrow \mathbb{N}$ by setting $L_{i+1}(v_k) = n_i + k$ and $L_{i+1}(u) = L_i(u)$ for all other $u \in \bar{\Theta}$. Then $w_r = [u_1, \dots, u_r]$. \square

We now give an upper bound for the growth rate of $\#W_R$, which will complete the proof of Theorem 2.1.

Lemma 2.5. $\lim_{R \rightarrow \infty} \frac{\log \#W_R}{R} \leq \log(2c_0 + 2) + 1.$

Proof. Suppose $p(R)$ and $q(R)$ are two functions of R with $\lim_{R \rightarrow \infty} \frac{p(R)}{q(R)} = \frac{1}{\epsilon}$. Then, using Stirling's formula, $\log \binom{p}{q}$ is asymptotic to $pH(\epsilon)$ as $R \rightarrow \infty$, where

$$H(\epsilon) = \epsilon \log \frac{1}{\epsilon} + (1 - \epsilon) \log \frac{1}{1 - \epsilon}$$

is the binary entropy function. (See [Mac03, Ch. 1].)

For any $R \geq 1$ and $C \geq R$, by a simple counting argument, the set

$$W(R, C) = \{(x_1, \dots, x_R) : x_i \in \{1, \dots, C\}, x_1 \leq \dots \leq x_R\}$$

has cardinality $\#W(R, C) = \binom{C + R - 1}{R}$.

Let $C = C_R + 1 = n + c_0 R + 1$. We have:

$$\#W_R \leq 2^R \binom{n + R(c_0 + 1)}{R} \quad \text{and} \quad \lim_{R \rightarrow \infty} \frac{n + R(c_0 + 1)}{R} = c_0 + 1.$$

Therefore,

$$\begin{aligned} \lim_{R \rightarrow \infty} \frac{\log \#W_R}{R} &\leq \lim_{R \rightarrow \infty} \frac{\log 2^R \binom{n + R(c_0 + 1)}{R}}{R} = \lim_{R \rightarrow \infty} \frac{R \log 2 + (n + R(c_0 + 1)) H\left(\frac{1}{c_0 + 1}\right)}{R} \\ &= \log 2 + (c_0 + 1) H\left(\frac{1}{c_0 + 1}\right) = \log 2 + \log(c_0 + 1) + c_0 \log\left(1 + \frac{1}{c_0}\right) \\ &\leq \log(2c_0 + 2) + 1. \end{aligned} \quad \square$$

3 Evolving structures on G -Labeled graphs

A graph is oriented if each edge is oriented. For any edge e of an oriented graph, denote by $i(e)$ and $t(e)$ the initial and terminal vertex of e . If e is a *loop*, then $i(e) = t(e)$. The orientation of e induces an orientation on each half-edge of e : the half-edge e_l containing $i(e)$ is oriented so that $i(e_l) = i(e)$ ($t(e_l)$ is a point in the interior of e), and the half-edge e_r containing $t(e)$ is oriented so that $t(e_r) = t(e)$.

Given a group G , an oriented graph is G -labeled if each edge is labeled by an element of G . Two G -labeled graphs are *equivalent* if one is obtained from the other by reversing the orientation of some of the edges and relabeling those edges with the inverse words. This defines an equivalent relation on the set of G -labeled graphs. Fix n and let $\mathcal{G}_n(G)$ be the set of equivalent classes of trivalent G -labeled graphs of rank n . (Recall the rank of a graph is the rank of its fundamental group.) We now consider operations that derive from an element in $\mathcal{G}_n(G)$ another element in $\mathcal{G}_n(G)$.




Figure 1: Double split and loop split

Let $\Gamma \in \mathcal{G}_n(G)$. Let e be an edge with label g_e . There are two *types* of edges in Γ : loop or non-loop. First assume e is not a loop. Choose a half-edge \tilde{a} (not a half-edge of e) incident at $i(e)$, and let a be the edge associated to \tilde{a} with label g_a . Disconnect \tilde{a} from $i(e)$ and reattach it to $t(e)$, while changing the label $g_a \rightarrow g_a g_e$ if $t(\tilde{a}) = i(e)$, or $g_a \rightarrow \overline{g_e} g_a$ if $i(\tilde{a}) = i(e)$. We call this a *forward split* along e . A forward split along e is well defined: if we reverse the orientation of a and invert g_a , then the resulting graph is equivalent. Similarly, take a half-edge \tilde{b} incident at $t(e)$. A *backward split* along e is obtained by disconnecting \tilde{b} from $t(e)$ and reattaching it to $i(e)$, while changing the label $g_b \rightarrow g_e g_b$ if $i(\tilde{b}) = t(e)$, or $g_b \rightarrow g_b \overline{g_e}$ if $t(\tilde{b}) = i(e)$. This is again well defined. A *double split* along e (see Figure 1) is the composition of a forward and a backward split along e . The resulting graph from a double split is trivalent, unlike from a backward or forward split alone. If we reverse the orientation of e and invert g_e , then a forward split along e becomes a backward split along e and vice versa. Therefore, a double split is well-defined on the equivalent class of Γ .

For a loop e , let a be the edge connected to e with label g_a . A *forward split* along e changes the label $g_a \rightarrow g_a g_e$ if $t(a) = i(e)$, or $g_a \rightarrow \overline{g_e} g_a$ if $i(a) = i(e)$. A *backward split* changes the label $g_a \rightarrow g_a \overline{g_e}$ if $t(a) = i(e)$, or $g_a \rightarrow g_e g_a$ if $i(a) = t(e)$ (see Figure 1). By a *loop split* along e we will mean either a forward or a backward split along e . This is again well defined on the equivalent class of Γ .

For any edge e of Γ , a *split* along e will mean either a double split or a loop split depending on the type of e . We will represent a split along e by $\Gamma \xrightarrow{s} \Gamma'$ and call $e = \text{supp}(s)$ the *support* of s .

Fix $\Gamma_0 \in \mathcal{G}_n(G)$. A derivation $D = [s_1, \dots, s_R]$ of length R is a sequence of splits

$$\Gamma_0 \xrightarrow{s_1} \Gamma_1 \xrightarrow{s_2} \dots \xrightarrow{s_R} \Gamma_R.$$

Set $\Gamma_i = [s_1, \dots, s_i](\Gamma_0)$; also write $\Gamma_R = D(\Gamma_0)$. We will say Γ_R is derived from Γ_0 by D . Let $B_R(\Gamma_0)$ be the set of all trivalent graphs (up to equivalence) that are derived from Γ_0 by a derivation of length at most R . Our main result is Theorem D, restated below.

Theorem 3.1. For any $\Gamma_0 \in \mathcal{G}_n(G)$,

$$\lim_{R \rightarrow \infty} \frac{\log \#B_R(\Gamma_0)}{R} \leq 3 \log 4.$$

The main idea behind the proof of Theorem 3.1 is to give a normal form for a derivation. This was done in [STT92] for unlabeled graphs. It turns out labeled graphs do not pose significant additional difficulties. We are also careful to obtain an explicit upper bound for the growth rate of $\#B_R(\Gamma)$.




Figure 2: Configurations of splits

A split $\Gamma \xrightarrow{s} \Gamma'$ defines a bijection between the edges of Γ and Γ' . For any edge e in Γ , let $s(e)$ be its image in Γ' . We say $\text{supp}(s)$ and its vertices are *destroyed* by s , and $s(\text{supp}(s))$ and its vertices are *created* by s . All other vertices of Γ *survive* s . An edge of Γ survives s if all of its vertices survive s .

Given $\Gamma \xrightarrow{s} \Gamma'$. If a representative of Γ is chosen, then s naturally induces representative for Γ' . Fix a representative in the equivalent class of Γ_0 . This way, for any derivation $D = [s_1, \dots, s_R]$, we can inductively define a representative for each $\Gamma_i = [s_1, \dots, s_i](\Gamma_0)$.

We will refer to Figure 2 for the following discussion. For each vertex v of Γ_0 , label the half-edges at v from 1 to 3 so they can be cyclically ordered. Let $D = [s_1, \dots, s_R]$ be a derivation. We will cyclically order the half-edges at each vertex of $\Gamma_i = [s_1, \dots, s_i](\Gamma_0)$ and *label* each s_i as follows. Let X be a fixed planar binary tree with four valence-1 vertices. The distinguished middle edge of X is oriented (see Figure 2). Let P be the planar graph which is the wedge of an interval and an oriented loop (also see Figure 2). Now suppose for $i \geq 0$, the half-edges of Γ_i are labeled. Let e be the support of s_{i+1} in Γ_i . If e is not a loop, then the cyclic ordering at $i(e)$ and $t(e)$ allows us to identify a contractible neighborhood of e with X . The four configurations in the left column of Figure 2 represent all possible double splits with support the middle edge of X . Record the label of the configuration that s_{i+1} identifies with; this is the label of s_{i+1} . Similarly, if e is a loop, then identify a neighborhood of e with P . Label s_{i+1} by 0 if s_{i+1} is a forward split along e and label s_{i+1} by 1 otherwise (see Figure 2). Let $\ell \in \{0, 1, 2, 3\}$ be the label of s_{i+1} . Note that we always know if e is a loop or not so there is no confusion with the duplication of the labels 0 and 1. For each vertex v of Γ_{i+1} , if v is not created by s_{i+1} , then the half-edges at v will inherit their labels from Γ_i ; otherwise, label the half-edges at v from 1 to 3 according to the right side of configuration ℓ in Figure 2.

Let $D = [s_1, \dots, s_R]$. Compute the label of each s_i from above. Fix $i \geq 0$ and let e be any edge in Γ_i . For $k > i + 1$, we will say e *survives* $[s_{i+1}, \dots, s_{k-1}]$ if for all $j = i + 1, \dots, k - 1$, the image of e in Γ_{j-1} survives s_j . In particular, e remains the same type from Γ_i to Γ_{k-1} . Let e_i be the preimage of $\text{supp}(s_k)$ in Γ_i . We say s_k is *ready* for Γ_i if e_i survives $[s_{i+1}, \dots, s_{k-1}]$. In this case, we can apply s_k to Γ_i with support e_i using the label of s_k ; this is well-defined since e_i is the same type as $\text{supp}(s_k)$.

Consider

$$\Gamma_{k-2} \xrightarrow{s_{k-1}} \Gamma_{k-1} \xrightarrow{s_k} \Gamma_k.$$

Suppose s_k is ready for Γ_{k-2} . Apply s_k to Γ_{k-2} and let Γ'_{k-1} be the resulting graph. Propagate the labels of half-edges from Γ_{k-2} to Γ'_{k-1} as before. Since e_{k-2} and $e = \text{supp}(s_{k-1})$ are disjoint in Γ_{k-2} , e survives s_k , so we may apply s_{k-1} to Γ'_{k-1} with support $s_k(e)$ using the label of s_{k-1} . Let

$$\Gamma_{k-2} \xrightarrow{s_k} \Gamma'_{k-1} \xrightarrow{s_{k-1}} \Gamma'_k$$

be the derivation obtained by switching the order of s_{k-1} and s_k . We claim the following:

Lemma 3.2. *With the same notation as above. If s_k is ready for Γ_{k-2} , then s_{k-1} and s_k commute; that is, $\Gamma_k = \Gamma'_k$.*




Figure 3: If s_k is ready for position $k - 2$, then s_{k-1} and s_k commute.

Proof. For any split s , we say an edge is *affected* by s if its label is changed by s . Any split affects at most two edges. Let e and e' be the supports of s_{k-1} and s_k in Γ_{k-1} respectively. If s_{k-1} and s_k do not affect the same edge in Γ_{k-2} , then they clearly commute. So let a be an edge in Γ_{k-2} affected by both s_{k-1} and s_k . a must share a vertex with both e and e' . Since e and e' are disjoint, a cannot be a loop. The proof that the labels of $s_{k-1} \circ s_k(a)$ and $s_k \circ s_{k-1}(a)$ are the same now follows from considering different cases. The proof in all cases are similar. Figure 3 shows the case when neither e and e' are loops and $t(e') = i(a)$ and $t(a) = i(e)$. Since $s_k \circ s_{k-1}$ and $s_{k-1} \circ s_k$ affect the edges labels the same way, $\Gamma'_k = \Gamma_k$. \square

Let $D = [s_1, \dots, s_R]$. For any $i \leq k - 1$, if s_k is ready for Γ_i , then s_k is ready for Γ_j , for all $i \leq j \leq k - 1$. By applying Lemma 3.2 $k - i$ times, we see that

$$D' = [s_1, \dots, s_i, s_k, s_{i+1}, \dots, \widehat{s_k}, \dots, s_R]$$

is a well defined derivation and $D(\Gamma_0) = D'(\Gamma_0)$.

Set $D_0 = D$ and let $\Gamma_R = D(\Gamma_0)$. Inductively, we will construct a sequence D_1, \dots, D_R of derivations such that $D_j(\Gamma_0) = \Gamma_R$ for all $j = 1, \dots, R$. The final derivation D_R will be called the *canonical derivation* of Γ_R coming from D_0 .

Let $M = 2n - 2$ be the number of vertices of Γ_0 . Label each vertex of Γ_0 by a distinct integer from 1 to M . Similarly, label the edges of Γ_0 from 1 to N , where $N = 3n - 3$ is the number of edges of Γ_0 .

Suppose for $i \geq 0$, $D_i = [u_1, \dots, u_i, t_{i+1}, \dots, t_R]$ has been constructed. Also, suppose the vertices and edges of Γ'_i have been labeled, where $\Gamma'_i = [u_1, \dots, u_i](\Gamma_0)$. Let U be the subset of $\{t_{i+1}, \dots, t_R\}$ consisting of splits that are ready for Γ'_i . Pick $t \in U$ such that t destroys the vertex of Γ'_i with the lowest label. Set

$$D_{i+1} = [u_1, \dots, u_i, t, t_{i+1}, \dots, \hat{t}, \dots, t_R].$$

Set $u_{i+1} = t$ and let $\Gamma'_{i+1} = [u_1, \dots, u_{i+1}](\Gamma_0)$. Let M_i and N_i be the maximal vertex and edge label of Γ'_i . Let f be the edge in Γ'_{i+1} created by t . Label f by $N_i + 1$. If f is a loop, then label its vertex by $M_i + 1$. If f is not a loop, then label $i(f)$ by $M_i + 1$ and $t(f)$ by $M_i + 2$. All other vertices and edges of Γ'_{i+1} will inherit their label from Γ'_i . This completes the construction.

Since Γ_0 has $2n - 2$ vertices and at most two vertices are created in each stage, the maximal vertex label of Γ_R is at most $2n - 2 + 2R$. Similarly, the maximal edge label of Γ_R is at most $3n - 3 + R$. Let

$$V = \{1, \dots, 2n - 2 + 2R\}, \quad E = \{1, \dots, 3n - 3 + R\}.$$

Let F_R be the set of all pairs of functions (ϕ, ψ) , where $\phi: V \rightarrow \{0, 1, 2, 3\}$ and $\psi: E \rightarrow \{0, 1, 2, 3\}$.

Proposition 3.3. *There exists an embedding of $B_R(\Gamma_0)$ into F_R , hence $\#B_R(\Gamma_0) \leq \#F_R$.*

Proof. By definition, any $\Gamma \in B_R(\Gamma_0)$ can be obtained from Γ_0 by a derivation D of length $r \leq R$. Let D_r be the canonical derivation of Γ coming from D . We now *encode* D_r by a pair of maps $\phi: V \rightarrow \{0, 1, 2, 3\}$ and $\psi: E \rightarrow \{0, 1, 2, 3\}$. Set D_r

$$\Gamma_0 \xrightarrow{u_1} \Gamma_1 \xrightarrow{u_2} \dots \xrightarrow{u_r} \Gamma_r.$$

For $i \in V$, let j be the largest index so that Γ_j has a vertex v with label i . If $j \geq R$ ($j > R$ means no such label exists), then set $\phi(i) = 0$. If $j < R$, then u_{j+1} must destroy v , so the support of u_{j+1} is an edge e in Γ_j where v is either $i(e)$ or $t(e)$. Define $\phi(i) \in \{1, 2, 3\}$ to be the label of the half-edge of e containing vertex v . If e is a loop, then choose $\phi(i)$ to be the label of any half-edge of e . For $i \in E$, let k be the largest index so that Γ_k has an edge e of label i . If $k \geq R$, then $\psi(i) = 0$. If $k < R$, then e is the support of t_{k+1} . In this case, define $\psi(i) \in \{0, \dots, 3\}$ to be the label of t_{k+1} .

To see this is an embedding, we will give a *decoding* procedure that will recover from (ϕ, ψ) the canonical derivation $D_r = [u_1, \dots, u_r]$ and hence Γ .

For each $k \geq 0$, suppose Γ'_k has been constructed. In Γ'_k , let i range from 1 to $2n - 2 + 2k$ in order and let v be the vertex in Γ_k with label i . If $\phi(i) = 0$, then move on to $i + 1$. Otherwise, $\phi(i)$ determines a unique edge e , where v is either the initial or terminal vertex of e , such that the half-edge of e at v has label $\phi(i)$. We now explain a *matching* procedure that can occur in two ways. If e is a loop, then we have a match. If e is not a loop, then let w be the other vertex of e with label i' . If $\phi(i')$ is exactly the label of the half-edge of e at w , then we have a match. In all other case, there is no match and we move on to $i + 1$. If there is a match, then let $j \in E$ be the label of e . The configuration

$\psi(j)$ determines a split supported on e which we call u'_{k+1} , and applying u'_{k+1} to Γ'_k yields Γ'_{k+1} . Proceed this way until $k = r$ results in a derivation D' .

To see that $D' = D_r$. Let e be the support of u_k . The encoding procedure ensures that the values of ϕ on the labels of $i(e)$ and $t(e)$ determine e , and hence a match, and the value of ψ on the label of e agrees with u_k . Furthermore, since only the splits that are ready at Γ_{k-1} can determine a match, and u_k is the unique one among them that destroys the vertex of Γ_{k-1} with the smallest label, the match coming from u_k will always be the first match the decoding procedure finds. This shows $\Gamma_k = \Gamma'_k$ and $t_k = t'_k$ for all k . Therefore, $B_R(\Gamma_0)$ embeds in F_R . \square

Since $\#F_R = 4^{5n-5+3R}$, $\lim_{R \rightarrow \infty} \frac{\log \#F_R}{R} = 3 \log 4$. This completes the proof of Theorem 3.1.

3.1 Triangulations of a surface

Let $\Sigma = \Sigma_{g,p}$ be an oriented surface of genus g with p punctures. For any $n \geq p$, let $\mathcal{T}_n = \mathcal{T}_n(\Sigma)$ be the set of homotopy classes of triangulations of Σ with n vertices (the punctures of Σ are always vertices of triangles.) A natural transformation of a triangulation is a *diagonal flip*. Given $T \in \mathcal{T}_n$. Let Δ and Δ' be two triangles in T that share a common edge E . View $\Delta \cup \Delta'$ as a quadrilateral with diagonal E . Replace E by the other diagonal in the quadrilateral yields a triangulation $T' \in \mathcal{T}_n$. Call this process a (diagonal) flip about E and denote it by $T \xrightarrow{d} T'$. Let $B_R(T)$ be the set of all triangulations of Σ obtained from T by a sequence of at most R diagonal flips.

Fix $T_0 \in \mathcal{T}_n$. Dual to T_0 is a trivalent graph Γ_0 obtained by putting a vertex in the interior of each triangle and connecting two vertices by an edge when two triangles share an edge. Pick a vertex x_0 in Γ_0 and let $G = \pi_1(\Sigma, x_0)$. We will label each edge of Γ_0 by an element of G as follows. Orient the edges of Γ_0 arbitrarily. Pick a spanning tree K_0 in Γ_0 and label each edge of K_0 by 1. Each edge e in the complement of K_0 represent an element of G : connect the end points of e to x_0 along K_0 and orient the resulting closed curve so that it matches the orientation of e in Γ_0 . Now, label e by the element that this closed curve represents in G . This makes a G -labeled graph $\Gamma_0 \in \mathcal{G}_m(G)$, where $m = 2g + n - 1$ is the rank of Γ_0 .

By a pair (Γ, f) we will mean a G -labeled graph $\Gamma \in \mathcal{G}_m(G)$ together with an embedding $f: \Gamma \rightarrow \Sigma$. We say a pair (Γ, f) is *well-labeled* if for any closed path p in Γ , the product of labels of edges along p is in the conjugacy class in G represented by $f(p)$. By construction, (Γ_0, i) , where Γ_0 is the dual graph to T_0 and i is the inclusion map, is well-labeled.

Proposition 3.4. *There exists an embedding of $B_R(T_0)$ into $B_R(\Gamma_0)$, hence $\#B_R(T_0) \leq \#B_R(\Gamma_0)$.*

Proof. Assume T and a well-labeled dual graph (Γ, i) are given. Consider a flip $T \xrightarrow{d} T'$ about an edge E in T and let e be the edge in Γ dual to E . Identify the quadrilateral containing E and a contractible neighborhood of e dual to the quadrilateral with the left-hand side of Figure 4. We define a split move $\Gamma \xrightarrow{s} \Gamma'$ supported on e , where (Γ', i) is also embedded in Σ , as indicated by Figure 4. We refer to s as the split associated to the flip d . Note that a closed path p in Γ can naturally be mapped to a homotopic closed path p' in Γ' and the products of labels along edges of p and p' are the same. That is, the pair (Γ', i) is still well-labeled.

We now define a map from $B_R(T_0)$ to $B_R(\Gamma_0)$. For any $T \in B_R(T_0)$, choose an arbitrary sequence of flips $T_0 \xrightarrow{d_1} T_1 \xrightarrow{d_2} \dots \xrightarrow{d_R} T_R = T$ and let $\Gamma_0 \xrightarrow{s_1} \Gamma_1 \xrightarrow{s_2} \dots \xrightarrow{s_R} \Gamma_R$ be the associated sequence of *dual*




Figure 4: Flip and dual split.

splits as constructed above. The map from $B_R(T_0)$ to $B_R(\Gamma_0)$ is defined by sending T_R to (Γ_R, i) and then to Γ_R .

We show that this map is injective. In fact, for triangulations T and T' and dual labeled graphs (Γ, i) and (Γ', i) that are well-labeled, we show that if Γ and Γ' are equivalent G -labeled graphs, then there exists a homeomorphism of Σ homotopic to the identity taking T to T' .

Since Γ and Γ' are equivalent, there is a graph isomorphism $\phi: \Gamma \rightarrow \Gamma'$ such that the label of any edge $e \in \Gamma$ matches the label of $\phi(e) \in \Gamma'$. Since Γ and Γ' are dual graphs to the triangulations T and T' respectively, we can build a homeomorphism $f: \Sigma \rightarrow \Sigma$ mapping a triangle of T associated to a vertex $v \in \Gamma$ to the triangle of T' associated to the vertex $\phi(v)$. To show that ϕ is homotopic to identity, it is sufficient to show that every closed path q in Σ is homotopic to $f(q)$.

First perturb q so it missed the vertices of T . Then q can be pushed to a closed path p in Γ . Since q is homotopic to p , we have $f(q)$ is homotopic to $p' = f(p)$. But the product of labels along the closed paths p and p' are identical, which means p and p' represent the same conjugacy class in G and hence are homotopic. This finishes the proof. \square

Theorem E from the introduction now follows from Theorem 3.1 and Proposition 3.4.

References

- [AP14] Jayadev S. Athreya and Amritanshu Prasad. Growth in right-angled groups and monoids. Preprint, 2014. Available at [arXiv:1409.4142](https://arxiv.org/abs/1409.4142) [math.GR].
- [Art47] E. Artin. Theory of braids. *Ann. of Math. (2)*, 48:101–126, 1947.
- [Bir74] Joan S. Birman. *Braids, links, and mapping class groups*. Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 82.
- [FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [GdlH97] R. Grigorchuk and P. de la Harpe. On problems related to growth, entropy, and spectrum in group theory. *J. Dynam. Control Systems*, 3(1):51–89, 1997.
- [HM95] Susan Hermiller and John Meier. Algorithms and geometry for graph products of groups. *J. Algebra*, 171(1):230–257, 1995.

- [Lic64] W. B. R. Lickorish. A finite set of generators for the homeotopy group of a 2-manifold. *Proc. Cambridge Philos. Soc.*, 60:769–778, 1964.
- [LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin-New York, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.
- [Mac03] David J. C. MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, New York, 2003.
- [McM14] Curtis McMullen. Entropy and the clique polynomial. *J. Topol.*, 7(29), 2014.
- [MKS66] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory: Presentations of groups in terms of generators and relations*. Interscience Publishers [John Wiley & Sons, Inc.], New York-London-Sydney, 1966.
- [RT13] Kasra Rafi and Jing Tao. The diameter of the thick part of moduli space and simultaneous Whitehead moves. *Duke Math. J.*, 162(10):1833–1876, 2013.
- [Sco07] Richard Scott. Growth series for vertex-regular $CAT(0)$ cube complexes. *Algebr. Geom. Topol.*, 7:285–300, 2007.
- [STT92] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Short encodings of evolving structures. *SIAM J. Discrete Math.*, 5(3):428–450, 1992.